

AMENDMENTS TO THE CLAIMS

1. (Currently Amended) A method for operating a storage volume, comprising:

establishing a profile for the storage volume, said profile being established using a predictive algorithm that is a function of at least one of:

a predicted read activity that will likely address the volume;

a predicted write activity that will likely address the volume; and

an amount of data likely to be read or written to the volume; and

while accepting and executing a host command for said storage volume:

collecting statistical data related to said storage volume, said statistical data including a total number of sectors read and a total number of sectors written, and independently maintaining said collected statistical data for each storage subvolume, each subvolume having a programmable granularity; and

analyzing said statistical data to determine whether a performance of said storage volume meets a goal;

if said goal is not met,

~~allocating or deleting a cluster;~~

modifying said profile, the statistical data and a current state of said profile to create a new state of said profile; and

modifying said storage volume to be consistent with the new state of said profile.

2. (Original) The method of claim 1, wherein said profile comprises information related to at least one of:

a size of the storage volume;

a RAID type identifier associated with clusters of the storage volume; and

a preferred location for clusters of the storage volume.

3. (Original) The method of claim 2, wherein said preferred location is specified as a function of a radius from a cluster to a center of a storage device.

4. Cancelled.

5. (Previously presented) The method of claim 1, wherein said predictive algorithm is also a function of a quality of service requirement for the storage volume.

6. (Previously presented) The method of claim 5, wherein said predictive algorithm is also a function of at least one of:

a redundancy requirement for the storage volume;

a size of the storage volume;

a number of storage devices in the storage volume; and

a performance expectation for the volume.

7. (Original) The method of claim 1, wherein said step of establishing a profile comprises storing said profile in a memory.

8. (Original) The method of claim 1, wherein said step of collecting statistical data comprises collecting at least one of a following statistical components:

a total number of read commands processed;

a total number of write commands processed;

a total number of times a command queue for holding pending read or write commands has queue length which coincides with one of a plurality of queue length ranges; and

a total number of times a latency for processing a read or write command coincides with one of a plurality of latency ranges.

9. Cancelled.

10. Cancelled.

11. (Original) The method of claim 1, wherein said step of collecting statistical data comprises forming an I/O density histogram.

12. (Original) The method of claim 1, wherein said step of modifying said profile comprises:

changing a RAID type of a cluster to optimize a frequently performed operation.

13. (Original) The method of claim 12, wherein said step of modifying said profile comprises detecting whether a cluster has significantly more write transactions than

read transactions and, if said cluster has significantly more write transactions than read transactions, changing a RAID type of the cluster from a type with redundancy through parity to a type with redundancy through mirroring.

14. (Original) The method of claim 12, wherein said step of modifying said profile comprises: detecting whether a cluster has significantly number of sequential data burst transfers and, if said cluster has significantly more data burst transfers, changing a RAID type of the cluster to a type with redundancy through parity.

15. (Original) The method of claim 12, wherein said step of modifying said profile comprises reallocating a cluster which is located at a relatively small radius from the center of a storage device and requires greater performance to a location having a greater radius from the center of the storage device.

16. (Original) The method of claim 1, wherein said goal is based on a quality of service requirement.

17. (Original) The method of claim 1, wherein said step of modifying said storage volume comprises determining whether a state of said storage volume is different from said profile and, if the state of said storage volume is different from said profile, determining whether an allocation status consistent with said profile is possible and, if said allocation status consistent with said profile is possible, reallocating clusters of said storage volume to become consistent with said profile.

18. (Original) The method of claim 1, wherein said host commands are received from a network.

19. (Currently Amended) A storage system, comprising:

a controller having a first interface for communicating with one or more hosts and a second interface; and

a plurality of storage devices, coupled to the controller via the second interface;

wherein said controller operates the plurality of storage devices to appear as one or more volumes to the one or more hosts, and for at least one volume the controller accepts a quality of service metric and enforces the quality of service metric by:

establishing a profile for the storage volume, said profile established as a function of at least one of:

a predicted read activity that will likely address the volume;

a predicted write activity that will likely address the volume; and

an amount of data likely to be read or written to the volume; and

while accepting and executing a host command for said storage volume,

collecting a statistical data related to said storage volume, said statistical data including a total number of sectors read and a total number of sectors written, and independently maintaining said collected statistical data for each storage subvolume, each subvolume having a programmable granularity; and

analyzing said statistical data to determine whether a performance of said storage volume meets a goal;

if said goal is not met,

~~allocating or deleting a cluster;~~

modifying said profile, the statistical data and a current state of said profile to create a new state of said profile; and

modifying said storage volume to be consistent with the new state of said profile.

20. (Currently Amended) A computer system, comprising:

at least one host;

a storage system, coupled to said at least one host, said storage system further comprising:

a controller; and

a plurality of storage devices, coupled to said controller;

wherein said controller operates the plurality of storage devices to appear as one or more volumes to the host, and for at least one volume the controller accepts a quality of service metric and enforces the quality of service metric by:

establishing a profile for the storage volume, said profile established as a function of at least one of:

a predicted read activity that will likely address the volume;

a predicted write activity that will likely address the volume; and

an amount of data likely to be read or written to the volume; and

while accepting and executing a host command for said storage volume,

collecting a statistical data related to said storage volume, said statistical data including a total number of sectors read and a total number of sectors written, and independently maintaining said collected statistical data for each storage subvolume, each subvolume having a programmable granularity; and

analyzing said statistical data to determine whether a performance of said storage volume meets a goal;

if said goal is not met,

~~allocating or deleting a cluster;~~

modifying said profile, the statistical data and a current state of said profile to create a new state of said profile; and

modifying said storage volume to be consistent with the new state of said profile.

21. (New) A method for operating a storage volume, comprising:

establishing a profile for the storage volume, said profile comprising information related to at least one of:

a size of the storage volume;

a RAID type identifier associated with clusters of the storage volume; and

a preferred location for clusters of the storage volume, said preferred location being specified as a function of a radius from a cluster to a center of a storage device;

said profile being established using a predictive algorithm that is a function of at least one of:

a predicted read activity that will likely address the volume;

a predicted write activity that will likely address the volume; and

an amount of data likely to be read or written to the volume; and

while accepting and executing a host command for said storage volume:

collecting statistical data related to said storage volume; and

analyzing said statistical data to determine whether a performance of said storage volume meets a goal;

if said goal is not met,

modifying said profile, the statistical data and a current state of said profile to create a new state of said profile; and

modifying said storage volume to be consistent with the new state of said profile.

22. (New) The method of claim 21, wherein said predictive algorithm is also a function of a quality of service requirement for the storage volume.

23. (New) The method of claim 22, wherein said predictive algorithm is also a function of at least one of:

a redundancy requirement for the storage volume;

a size of the storage volume;

a number of storage devices in the storage volume; and

a performance expectation for the volume.

24. (New) The method of claim 21, wherein said step of establishing a profile comprises storing said profile in a memory.

25. (New) The method of claim 21, wherein said step of collecting statistical data comprises collecting at least one of a following statistical components:

a total number of read commands processed;

a total number of write commands processed;

a total number of times a command queue for holding pending read or write commands has queue length which coincides with one of a plurality of queue length ranges; and

a total number of times a latency for processing a read or write command coincides with one of a plurality of latency ranges.

26. (New) The method of claim 21, wherein said step of collecting statistical data comprises forming an I/O density histogram.

27. (New) The method of claim 21, wherein said step of modifying said profile comprises:

changing a RAID type of a cluster to optimize a frequently performed operation.

28. (New) The method of claim 27, wherein said step of modifying said profile comprises detecting whether a cluster has significantly more write transactions than read transactions and, if said cluster has significantly more write transactions than read transactions, changing a RAID type of the cluster from a type with redundancy through parity to a type with redundancy through mirroring.

29. (New) The method of claim 27, wherein said step of modifying said profile comprises detecting whether a cluster has significantly number of sequential data burst transfers and, if said cluster has significantly more data burst transfers, changing a RAID type of the cluster to a type with redundancy through parity.

30. (New) The method of claim 27, wherein said step of modifying said profile comprises reallocating a cluster which is located at a relatively small radius from the center of a storage device and requires greater performance to a location having a greater radius from the center of the storage device.

31. (New) The method of claim 21, wherein said goal is based on a quality of service requirement.

32. (New) The method of claim 21, wherein said step of modifying said storage volume comprises determining whether a state of said storage volume is different from said profile and, if the state of said storage volume is different form said profile, determining whether an allocation status consistent with said profile is possible and, if said allocation status consistent with said profile is possible, reallocating clusters of said storage volume to become consistent with said profile.